



12 **EUROPEAN PATENT APPLICATION**

21 Application number : **95301758.9**

51 Int. Cl.⁶ : **H04Q 11/04, H04L 12/56**

22 Date of filing : **16.03.95**

30 Priority : **25.03.94 US 218015**

43 Date of publication of application :
27.09.95 Bulletin 95/39

84 Designated Contracting States :
DE ES FR GB IT

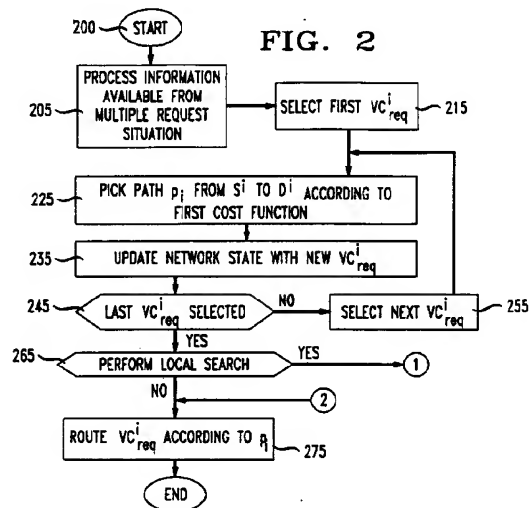
71 Applicant : **AT & T Corp.**
32 Avenue of the Americas
New York, NY 10013-2412 (US)

72 Inventor : **Gawlick, Rainer**
103 Pinckneystreet,
Apt.4
Boston, Massachusetts 02139 (US)
Inventor : **Kalmanek, Charles Robert, Jr.**
86 Great Hills Road
Short Hills, New Jersey 07078 (US)
Inventor : **Ramakrishnan, Kajamalai**
Gopalaswamy
146 Maple Avenue
Berkeley Heights, New Jersey 07922 (US)

74 Representative : **Watts, Christopher Malcolm**
Kelway, Dr. et al
AT&T (UK) Ltd.
5, Mornington Road
Woodford Green Essex, IG8 0TU (GB)

54 **A method of on-line permanent virtual circuit routing.**

57 A method of on-line routing of permanent virtual circuits is disclosed. The method uses an exponential cost function based on the network state and virtual circuit parameters to determine paths for the virtual circuits without scaling. The method may also advantageously use information available when multiple virtual circuits are requested to determine paths for the virtual circuits. Additionally, the method allows the routing selection to be refined so that the total cost of routing all the virtual circuits is reduced.



Technical Field

The invention relates to the routing of permanent virtual circuits in networks.

Background of the Invention

Computer networks are a principal means of exchanging or transferring information (e.g. data, voice, text, video, etc.) among host machines connected to the network. The network comprises nodes connected, to each other and to the hosts, by links. Typically, each link is bidirectional, i.e. information may be conveyed in forward and reverse directions, and each link is characterized by a bandwidth capacity in each direction.

An important consideration in network operation is how the information is routed. When information is to be exchanged between two particular hosts, a bidirectional path is established in the network between them. Typically, this path that is established is a so-called "virtual circuit" (VC), by which it is meant that a host simply specifies the destination for the information, and the network delivers the information as though a circuit is connected to the destination. One of many different routes and techniques could be selected to deliver the information, but the particular selection is of no concern to the host. The task of routing is to select the nodes and links between the nodes that comprise the path taken by the VC so as to efficiently utilize network resources, e.g. route as many VCs as possible without exceeding the bandwidth capacity of any particular link. This is often achieved by selecting a path so as to minimize some cost function that reflects the amount of network resources, e.g. bandwidth of links in the path, required by the selected path. Although a variety of cost functions can be used, cost functions typically take into account the current network state (i.e., the network topology and current allocation and availability of network resources), delay through the network, etc. The routing problem is often further complicated in that the routing must be effectuated "on-line," i.e. without knowledge of what future routing demands will have on network resources. Although this problem may be solved by so-called "dynamic-rerouting" techniques, these techniques typically adversely affect the quality of service offered to users of the network.

The manner in which the routing problem for virtual circuits is addressed typically is based on whether the VCs are permanent or switched and on whether the routing is done with incomplete or complete information. Permanent VCs are paths for information transfer between hosts designed to operate and remain established for long periods of time, perhaps on the order of years. Switched VCs are designed to operate for hours or days, and thus networks of switched VCs are reconfigured more often. See, B. Awerbuch et al., "Throughput-Competitive On-line Routing," *34th Annual Symp. on Foundations of Comp. Sci.*, Palo Alto, CA, November 1993. Routing with incomplete information means simply that the state of the network is not known or, if the state is known, that the state information is dated, e.g., that the available state information does not incorporate information regarding network resources allocated to the most recently routed VCs. Conversely, routing with complete information means that the state information is fully known and that the state information is up-to-date.

Previous techniques for routing permanent VCs have employed minimum hop routing in which the path going through the shortest number of nodes is selected. Recently, another method, employing exponential cost functions and scaling, has been suggested. See, J. Aspnes et al., "On-Line Load Balancing with Applications to Machine Scheduling and Virtual Circuit Routing," *Proc. 23rd Annual Symp. on Theory of Computing*, San Diego, CA, May 1993. In the scaling technique, a portion γ of the bandwidth capacity of each link is initially apportioned, and a cost function is computed for routing a path given that apportioned bandwidth. When routing can no longer be achieved in the network with that apportioned bandwidth, more bandwidth can be apportioned, i.e. the scale factor γ is increased. Typically, the function for determining the cost for a given link in a path for the requested VC is $C_l(x, \Delta x) = \alpha^{\gamma x_l + \gamma \Delta x_l} - \alpha^{\gamma x_l}$ where $C_l(x, \Delta x)$ is the cost for link l in the path, α is a constant, x_l is the fraction of the bandwidth capacity of the link that is in use and Δx_l is the fraction of the bandwidth capacity of the link that is requested by the VC. Note that this is an exponential function in that the constant α is raised to a power - the power being a function of γ and Δx_l and of γ and x_l . Thus, the cost function includes a term in which the network resource (i.e. bandwidth) is in scaled form in the exponent of the term.

These prior techniques of permanent VC routing have shortcomings. For example, minimum hop techniques do not adequately reflect the current state of the network in path selection. Additionally, the scaling method, by restricting itself to a fraction of the total link bandwidth capacity, may select costly paths even when demand for VCs is low.

Summary of the Invention

Central to the invention is a recognition that permanent virtual circuits may be routed on-line using a) an

exponential cost function and b) without scaling, and accordingly the invention is directed to a method of routing of permanent virtual circuits on these bases. The inventive method routes permanent virtual circuits in a network characterized by a network state by receiving a set of requests. Each request in the set is a request to route a permanent virtual circuit on a path, and each request is specified by one or more parameters. Each request is then routed on a selected path through the network according to an exponential cost function without scaling. The selected path is chosen to satisfy the parameters specifying the request, and the exponential function is a function of the network state and of one or more parameters specifying the request.

Brief Description of the Drawings

Other features and advantages of the invention will become apparent from the following detailed description taken together with the drawings in which:

FIG. 1 illustrates a centralized routing network in which the invention may be practiced.

FIG. 2 is a flowchart of the steps in the inventive method.

FIG. 3 is a flowchart of steps of a local search feature of the inventive method.

FIG. 4 illustrates a distributed routing network in which the invention may be practiced.

Detailed Description

FIG. 1 illustrates the structure of a network in which the invention may be practiced. Hosts $102-i, i = 1, 2, \dots$, exchange information via network 106. Network 106 comprises links $110-k, k = 1, 2$ connecting nodes $108-j, j = 1, 2, \dots$, to each other and to hosts $102-i$. A pair of nodes may be connected by one or more links.

Network 106 in FIG. 1 is a centralized routing system in that network 106 utilizes complete state information through use of centralized routing request processor 111. Request processor 111 is connected to hosts $102-i, i = 1, 2, \dots$ and to all nodes $108-j, j = 1, 2, \dots$. Request processor 111 has complete information about the network state. Thus, the cost for any path (i.e. the additional network resources required for any path) through the network can be determined, and, using the inventive method described below, all VCs in network 106 of FIG. 1 can be routed efficiently with respect to a given criterion, e.g. maximizing the amount of bandwidth in use among links.

FIG. 2 is a flowchart of the steps in the inventive method in which permanent VCs are routed on-line using an exponential cost function. The method of FIG. 2 routes permanent VCs using an exponential cost function without scaling. Since the first cost function may advantageously be based on a parameter such as the total bandwidth capacity among potential links in a path, scaling is eliminated.

In step 200 of FIG. 2 at any given time, the network must respond to set of i_{\max} , $i_{\max} = 0, 1, 2, \dots$, requests for establishing VCs. Each individual request, $VC_{req}^i, i = 1, \dots, i_{\max}$ is specified by one or more parameters. For example, each VC_{req}^i may be specified by the source host S^i , destination host D^i , bandwidth requested in the forward direction B_f^i and bandwidth requested in the reverse direction B_r^i . Thus,

$$VC_{req}^i = (S^i, D^i, B_f^i, B_r^i)$$

Step 205 of FIG. 2 also illustrates an optional feature that may be used with the inventive method. The optional feature, disclosed in co-pending application "A Method of Routing Multiple Virtual Circuits," filed concurrently herewith and commonly assigned, advantageously uses information available in the multiple VC request situation to route the VCs. In step 205 of FIG. 2 each request in the set of requests is routed as a function of one or more parameters of a plurality of requests in the set. For example, the requests can be routed on a first-come, first-served basis but with knowledge of what the average bandwidth requirement is of the requests in the set of requests. Thus, when a particular request is routed, the cost function for the routing may reflect whether the request requires a large or small amount of bandwidth relative to other requests in the set thereby causing relatively small bandwidth requests to be routed on links already near capacity so as to conserve bandwidth on other links for routing large bandwidth requests. As another example, the requests in the set of requests may be optionally ordered with respect to one or more of the parameters. The ordering process enables the inventive method to route first those requests requiring the most network resources according to the cost function and thus requiring the most flexibility in routing. In particular, if an objective in the routing is to conserve bandwidth, it is advantageous to order VC_{req}^i in decreasing order of total requested bandwidth (forward and reverse) so that those requests requiring large bandwidths can be accommodated without increasing the risk of exceeding the bandwidth capacity of any link. Other orderings are possible depending on network performance goals.

Returning to FIG. 2 the first request is selected in step 215, and a path, p_i , from S^i to D^i for the request is selected based on a first cost function. Note - the term "first cost function" is used here to determine a path

for each VC_{req}^i and is to be distinguished from a "second cost function" used in the local search described below.

First cost function $Cost^1$ is advantageously chosen to be a measure of the cost of a particular path over a set of links in the particular path for a given VC. The one particular path with the minimum cost is then selected as the path associated with the given VC. In the inventive method it is recognized that $Cost^1$ can be advantageously chosen to be an exponential function. Based on the specification of the VCs as above, the function is selected to be an exponential function of four variables $x_{f,h}^i$, $\Delta B_{f,h}^i$, $x_{r,l}^i$ and $\Delta B_{r,l}^i$ where $x_{f,h}^i$ is the percentage of the bandwidth capacity currently in use on the link in the forward direction, and $x_{r,l}^i$ is the fraction of the bandwidth capacity currently in use in the reverse direction. Parameters $\Delta B_{f,l}^i$ and $\Delta B_{r,l}^i$ are the fraction of the bandwidth capacity requested in the forward and reverse directions, respectively. For example, $Cost^1$ may advantageously be chosen as

$$Cost^1 = \sum_{\text{links } l \text{ in the path}} \left[(A)^{x_{f,l}^i + \Delta B_{f,l}^i} - (A)^{x_{f,l}^i} \right] + C \Delta B_{f,l}^i + \sum_{\text{links } l \text{ in the path}} \left[(A)^{x_{r,l}^i + \Delta B_{r,l}^i} - (A)^{x_{r,l}^i} \right] + C \Delta B_{r,l}^i \quad (1)$$

where A and C are selected constants. The constant A is related to how closely the method follows a minimum hop method - the smaller the value of A the closer to the minimum hop technique. The constant C is used dampen the effect of the value of A when little bandwidth is in use in a network. Illustratively, the constant A may be selected as 100 - indicating that the method does not closely follow the minimum hop method. Constant C is typically selected from the range $0 \leq C \leq 1$, with C advantageously close to zero indicating a small amount of damping. Thus, first cost function is an exponential function with constant A raised to powers that are functions of $x_{f,h}^i$, $\Delta B_{f,h}^i$, $x_{r,l}^i$ and $\Delta B_{r,l}^i$. Note that the function does not include a scale factor. Instead, first cost function is an exponential function which includes a term having an exponent where the exponent includes the allocated amount of a network resource (in this case, bandwidth) in unscaled form. Thus, unlike exponential functions which use a scale factor to reduce the bandwidth capacity of a link to some apportioned amount, first cost function may be used to select paths based on the available bandwidth in each link.

In strict mathematical terms, first cost function may advantageously be characterized as, over the relevant domain, having both a positive semi-definite Hessian and a non-negative gradient in each direction. For the four normalized variables above, $x_{f,h}^i$, $\Delta B_{f,h}^i$, $x_{r,l}^i$ and $\Delta B_{r,l}^i$, the relevant domain for first cost function is $[0,1] \times [0,1] \times [0,1] \times [0,1]$ since each variable is a fraction of the capacity of a link. However, those skilled in the art will recognize that functions equivalent to (1) can be formulated. For example, the variables need not be normalized with respect to link capacity if the domain is properly adjusted or an exponential cost function with a negative semi-definite Hessian and with a non-positive gradient in each direction can be used as long as the highest cost path based on that cost function is selected.

Once the minimum cost path for the first request has been selected, the network state is updated to reflect the resources allocated for the first request in step 235. Steps 225 and 235 are repeated in step 255 for each succeeding VC_{req}^i until all requests in the set have been filled.

Thus, the method of FIG. 2 routes permanent VCs using an exponential cost function without scaling. Specifically, since the exponential cost function is based on the total bandwidth capacity among potential links in a path, scaling is eliminated. The routing of VCs according to steps 215-255 permits "on-line" routing of multiple requests. The method also allows routing of those requests using information about the parameters of a plurality of concurrent requests.

Once all the VC_{req}^i have been addressed (step 245), an optional local search as disclosed in the co-pending application, *supra*, can be performed to refine the selection of paths so that the total cost of routing all VCs can be reduced. The steps in one embodiment of the local search of step 265 are illustrated in the flowchart of FIG. 3. In step 305 variables are initialized. In particular, P is defined as the set of paths $\{p_1, p_2, \dots, p_{i_{max}}\}$ associated with the i_{max} VC requests as determined by a routing method, e.g. the method of FIG. 2. The current set of best paths for routing is stored in P^* and the cost of routing that current set of best paths is C^* . Initially, $P^* = P$ and $C^* = C(P)$ where C is a second cost function, described below, which determines the cost of routing all the VC_{req}^i . In step 310 a flag, called *impflag*, is set to zero and a counter $i, i = 1, 2 \dots i_{max}$, is set to 1.

In step 315 a particular VC_{req}^i and its associated path p_i are selected. For illustrative purposes in FIG. 3 the VC_{req}^i are selected in increasing order of i . For the selected i a search is made for a lowest cost path p^* for

routing VC_{req}^i (using, for example, $Cost^1$) assuming all other VC_{req}^i are routed as in the set P . A new set of paths, \bar{P} , is formed in step 320 by setting \bar{P} to P except for the selected p_i which is set to p^* . In step 330 if $C(\bar{P}) < C^*$, then P^* and C^* are set to \bar{P} and $C(\bar{P})$, respectively, and $impflag$ is set to one.

Steps 315-340 are repeated for each successive VC_{req}^i . When all $i, i = 1, 2, \dots, i_{max}$, have been examined for potential rerouting, $impflag$ is checked to see if any improvement in routing is possible. If $impflag$ is zero, no alternative routing of any one VC_{req}^i will reduce the cost of routing all VC_{req}^i and the search is ended. As indicated in step 345, if $impflag$ has been set to one, there exists an alternative routing (i.e. a new shortest path) for one VC_{req}^i that results in the greatest reduction in routing cost. In step 355, that new lowest cost path, reflected in P^* , becomes P . The VC_{req}^i are now to be routed according to P^* at cost $C(P^*)$. Steps 310-355 are repeated until no new paths are found that reduce the cost.

The second cost function may advantageously be selected as

$$\begin{aligned}
 Cost^2 = & \sum_{\text{all links } l} \left[\left[(A)^{x_{f,l}^i} + \sum \Delta B_{f,l}^i - (A)^{x_{r,l}^i} \right] + C \Delta B_{f,l}^i \right] \\
 & + \sum_{\text{all links } l} \left[\left[(A)^{x_{r,l}^i} + \sum \Delta B_{r,l}^i - (A)^{x_{f,l}^i} \right] + C \Delta B_{r,l}^i \right] \quad (2)
 \end{aligned}$$

where the variables in the equation are defined as for equation 1 and where the sum in each exponential and in the linear term is over all VC_{req}^i such that their associated path uses the link in the outer summation. Thus, the second cost function gives the total cost of routing multiple VC_{req}^i on their associated paths.

The steps in FIG. 3 reflect a so-called "greedy" heuristic in which each possible alternative path for each VC_{req}^i is examined to see which alternative path, if any, reduces the cost of routing by the greatest amount. That alternative path is then selected as the new path associated with the VC. The process is repeated until no alternative path for any VC_{req}^i will reduce the cost of routing all the requests. Those skilled in the art will recognize that other local searches utilizing other search criteria can be used. For example, instead of searching through all the alternative paths to find the one that reduces the cost of routing all VC_{req}^i the most, it may be sufficient simply to find a first alternative path that reduces costs by any amount. Then the search could simply go onto the next VC_{req}^i . This "less greedy" heuristic searches the solution space differently and can potentially converge to a better "local minimum" than the "greedy" heuristic. This is highly dependent on the list of VC_{req}^i being routed and the current state of the network.

FIG. 4 illustrates the structure of a distributed network 416, comprising nodes 418- m and links 420- n , in which the inventive method may be practiced. Network 416 is a distributed routing system in that each node 418- m periodically exchanges state information. The state information reflects the amount of network resources available or in use on a link from a node to every neighboring node. Thus, the state information can be used to determine the cost for any path through the network. However, unless the state information propagates quickly relative to the speed with which VCs are established and torn down, the information will be incomplete (e.g. dated). Thus, each node may have a different description of the network state, and this description is called the local network state. The methods described above can be used in a distributed routing system except that incomplete state information (i.e. the local network state) is used. Note however, the paths determined based on the incomplete information may no longer be available (e.g. the capacity of a link may have become exhausted since the most recent state information was received) when the VC is actually to be routed (as for example in the routing of step 275 of FIG. 2). Requests for VCs that are not successfully routed, e.g. those requests for which the selected path is not available, may be included in the next set of requests or may form a new set of requests. Note further that it is advantageous when routing in a distributed routing system not to update the local network state with state information received from other nodes until an entire set of requests in the multiple VC request situation has been routed. This ensures that procedures, such as the search routine of FIG. 3, will terminate.

This disclosure describes a method of on-line permanent virtual circuit routing. The method disclosed herein have been described without reference to specific hardware or software. Instead, the method has been described in such a manner that those skilled in the art can readily adapt such hardware and software as may be available or preferable for particular application.

Claims

1. A method of routing requests for permanent virtual circuits in a network having resources, said network characterized by a network state at least reflecting an allocated amount of a predetermined network resource, said method comprising the steps of:
 - receiving a set of requests, wherein each request in said set of requests is a request to route a permanent virtual circuit on a path and wherein each request is specified by one or more parameters; and for each request:
 - routing the request on a selected path through said network according to an exponential cost function, wherein said selected path satisfies said parameters specifying the request, wherein said exponential function is a function of the network state and of one or more of said parameters specifying the request and wherein said exponential function includes a term having an exponent, said exponent including said allocated amount of said predetermined network resource in unscaled form.
2. A method of routing requests for permanent virtual circuits in a network having resources, said network characterized by a network state reflecting an allocated amount of a predetermined network resource, said method comprising the steps of:
 - receiving a set of requests, wherein each request in said set of requests is a request to route a permanent virtual circuit on a path and wherein each request is specified by one or more parameters; and
 - routing each request on a selected path through said network according to an exponential cost function, wherein said exponential function is a function of the network state and of one or more of said parameters specifying each request and wherein said exponential function includes a term having an exponent, said exponent including said allocated amount of said predetermined network resource in unscaled form.
3. The method of any of the preceding claims further comprising the steps of:
 - selecting an alternative path to one of said selected paths,
 - determining the value of a second function using said alternative path, and
 - determining the value of the second function using said selected path, and if said alternative path improves the value of said second function relative to said selected path, then replacing said selected path with said alternative path.
4. The method of claim 3 further comprising the step of:
 - repeating the steps of selecting an alternative path, determining values for the second function and replacing until selecting no alternative path will improve the second function.
5. The method of claim 3 or claim 4 wherein said second function is a function of the network state and of one or more of said parameters.
6. The method of any of the preceding claims further comprising the step of establishing the requests in said set of requests in an order according to one or more of said parameters prior to the routing step, the routing step routing said requests in said order.
7. A method of routing a request for a permanent virtual circuit on a path in a network having resources, wherein said request is specified by one or more parameters and wherein said network is characterized by a network state reflecting an allocated amount of a predetermined network resource, said method comprising the steps of:
 - routing said request on a selected path through said network according to an exponential cost function, wherein said selected path satisfies said parameters specifying said request, wherein said exponential function is a function of the network state and of one or more of said parameters specifying said request and wherein said exponential function includes a term having an exponent, said exponent including said allocated amount of said predetermined network resource in unscaled form.
8. The method of any of the preceding claims further comprising the step of updating said network state to reflect said selected path.
9. The method of any of the preceding claims wherein the Hessian of said exponential function is positive

semi-definite and wherein the gradient of said exponential function is non-negative.

10. The method of any of the preceding claims wherein the routing step is carried out based on complete information about said network state.

5

11. The method of any of claims 1 to 9 wherein the routing step is carried out based on incomplete information about said network state.

12. The method of any of the preceding claims wherein said network comprises a set of links and wherein each selected path through said network is specified by one or more links in said set of links.

10

13. The method of claim 12 wherein each selected path connects a first host and a second host.

14. The method of any of the preceding claims wherein one of said parameters is bandwidth.

15

20

25

30

35

40

45

50

55

FIG. 1

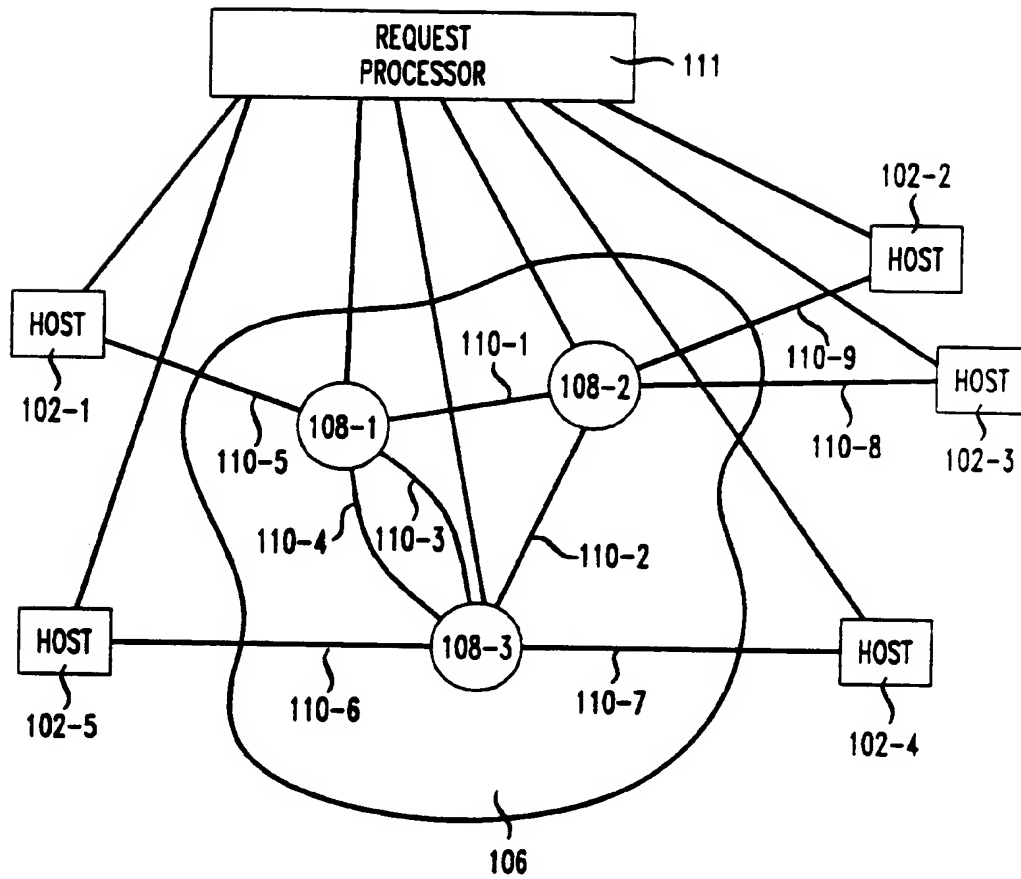


FIG. 2

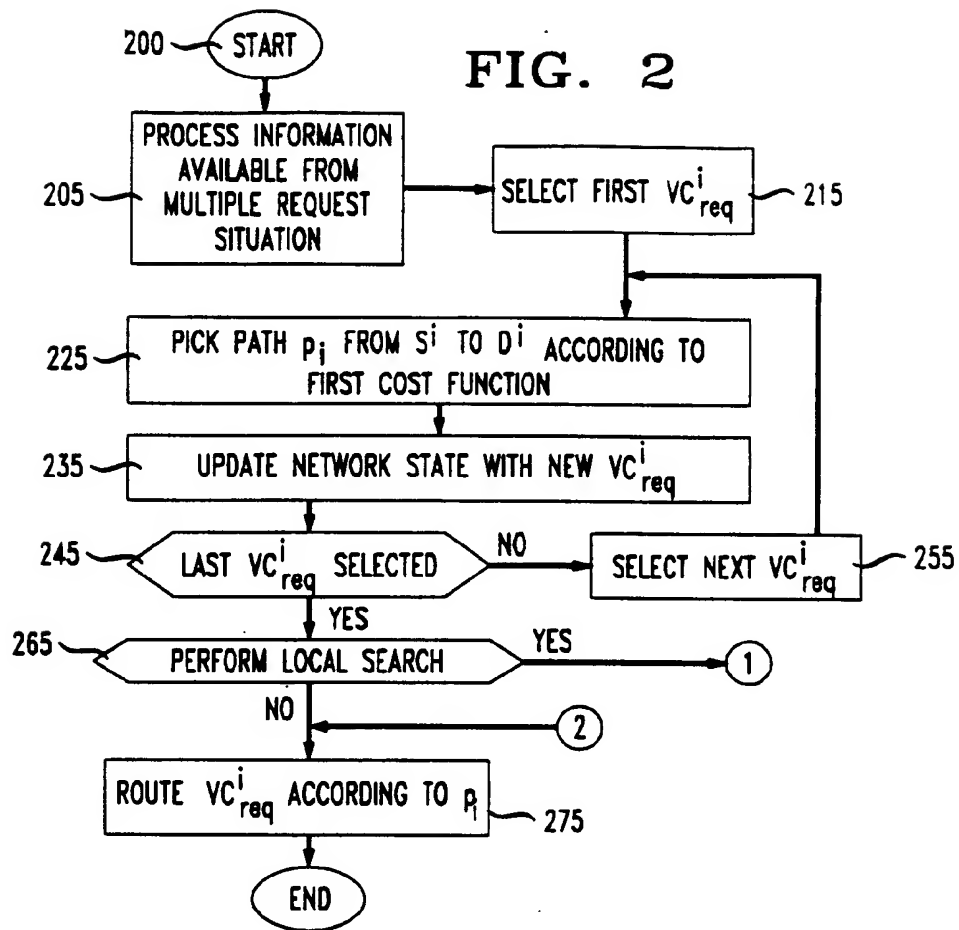


FIG. 4




FIG. 3

